

Partitioned System-Defined Aggregate Tables

AtScale Model Architects provide prioritized partitioning hints to the AtScale Engine to inform the creation of partitioned aggregate tables. The actual partitioning scheme used by the Engine depends on a number of factors, including:

- ▶ Whether the aggregate includes a column that matches a partition hint.
- ▶ Whether AtScale statistics suggest that partitioning would be worthwhile.
- ▶ Whether the target data warehouse supports table partitioning.

Defining Aggregate Partitions For Models

To define an aggregate partition scheme for a model, you must edit its underlying SML to include the `partitions` property. For more information, see [Models](#).

What To Do Next

After making your changes to the model's SML, you need to deploy the model so that it can be queried.

Once a query containing the dimension levels that you've specified as a partition has been executed on the deployed model, the AtScale engine determines whether to build a new system aggregate with the partition, based on the following criteria:

- ▶ If the estimated number of rows per partition meets or exceeds the threshold set in the engine setting `AGGREGATES.CREATE.PARTITION.SYSTEMDEFINEDAGGREGATE.THRESHOLD`.
- ▶ If the estimated number of partitions in the aggregate table exceeds the threshold set in `tables.create.partitions.maximumEstimatedNumberOfPartitions` the table will not be partitioned. The default value of this setting is 800.

Navigate to **Definitions** tab on the **Aggregates** page to see whether the system aggregate was created as a result of your query. For more information, see [Monitoring Aggregate Usage](#).